

# The PWRficient Processor

Peter Bannon  
VP Architecture

January 2006



# Outline

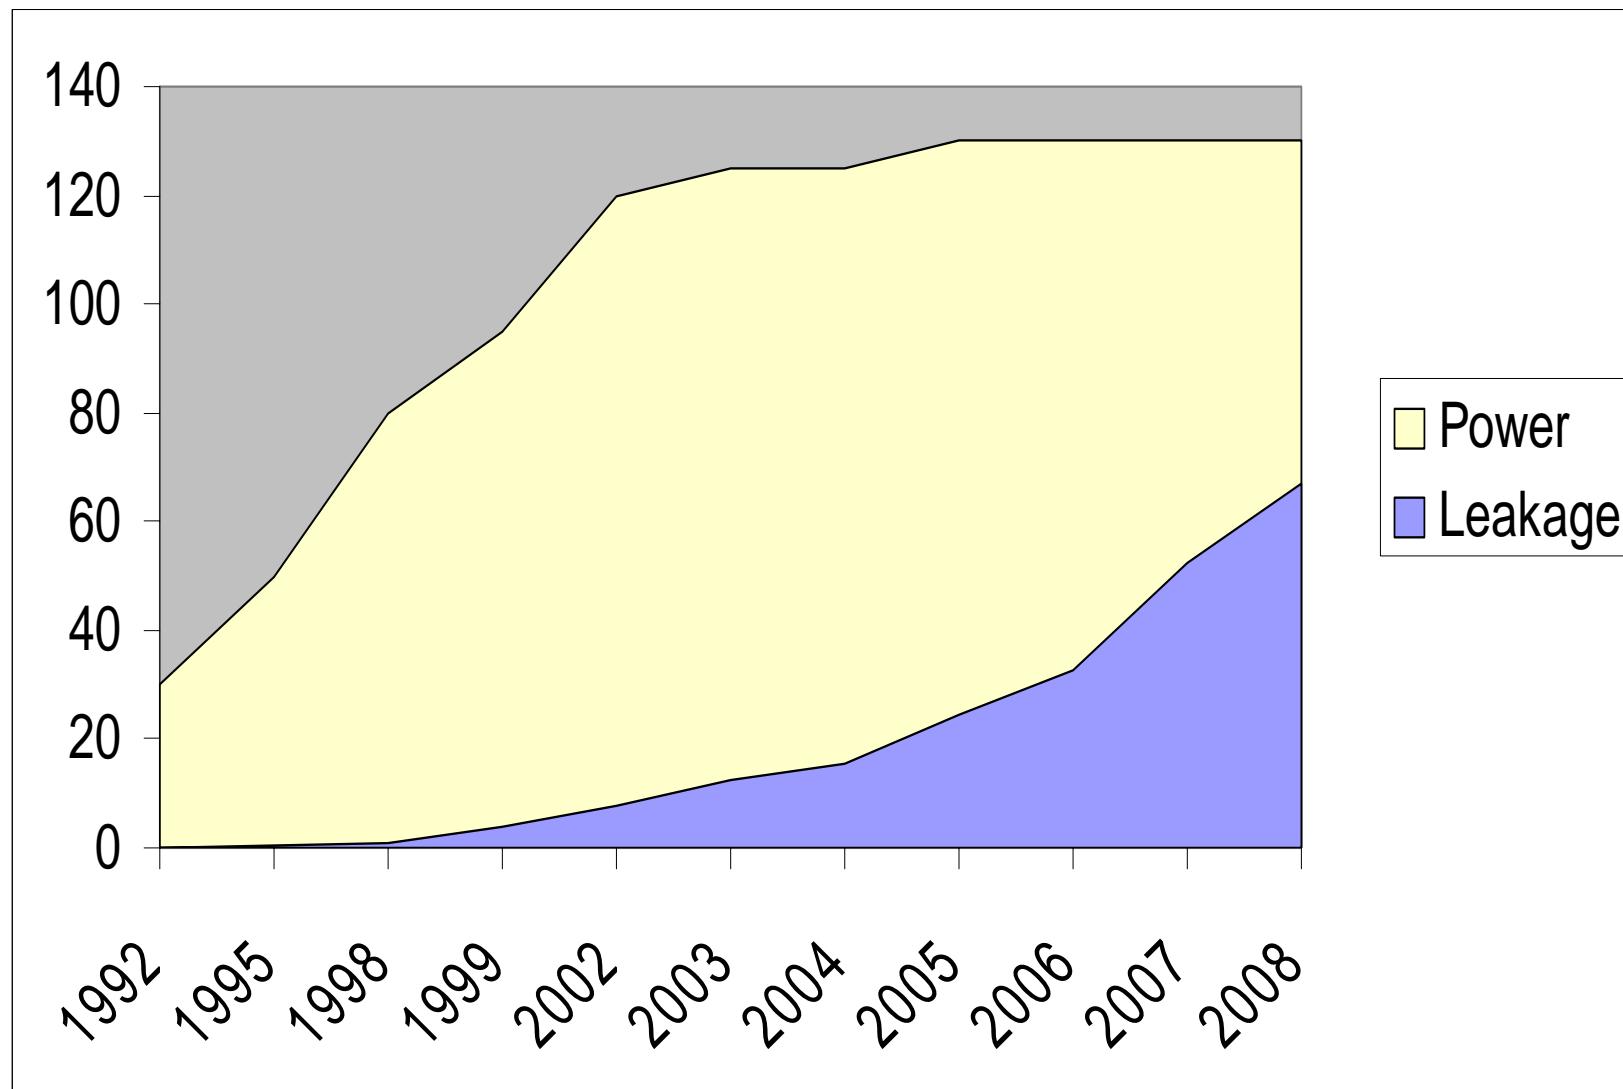
- ▶ PA Semi
- ▶ Low Power Design
- ▶ PWRficient 1682M platform processor
  - ▷ PA6T core
  - ▷ DDR2 memory
  - ▷ ENVOI I/O System
- ▶ Performance
- ▶ Power
- ▶ Conclusion

# PA Semi Overview

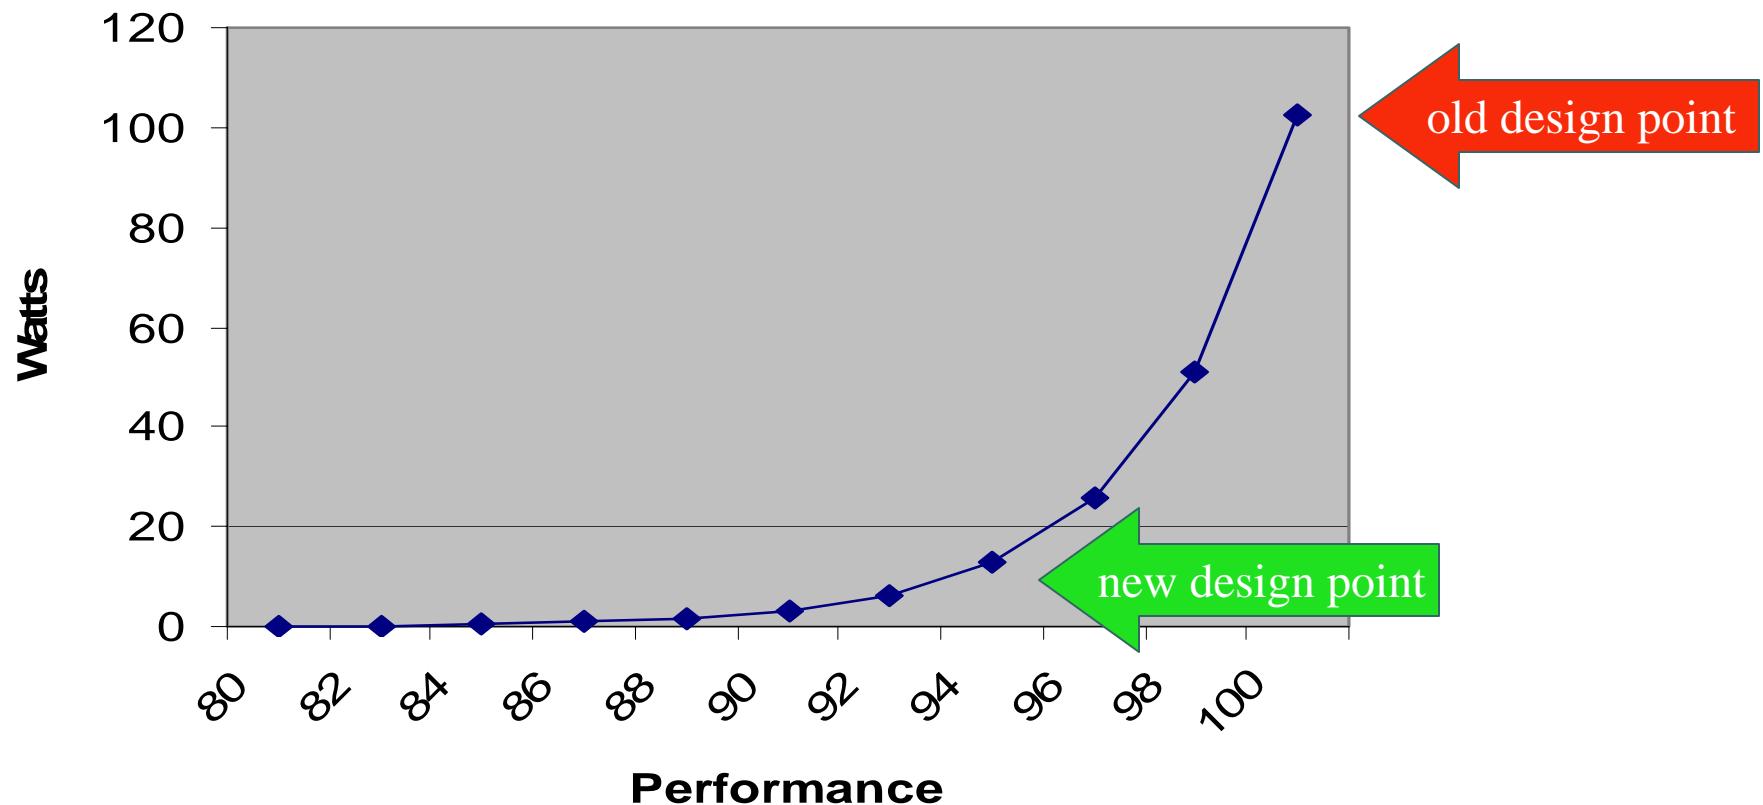
- ▶ P.A. Semi is a fabless semiconductor company located in Santa Clara, California specializing in high-performance, low-power microprocessor platforms
- ▶ Customers will select our products due to their high performance, low power, or low latency or the unique combination of those characteristics:
  - ▶ Embedded network, storage and telecom
  - ▶ Blade, grid and rack servers
  - ▶ Compute-intensive appliances

# Low Power Design

# Useful Power Is Decreasing



# Choosing Power Over Ultimate Performance



- ▶ Look for the exponential opportunities in power/performance
- ▶ Give up some performance for substantial power decrease

# Tuning the Design

## ► Tuning

- ▷ CMOS process target
- ▷ Circuit design style and sizing
- ▷ Micro-architecture features

## ► Integration

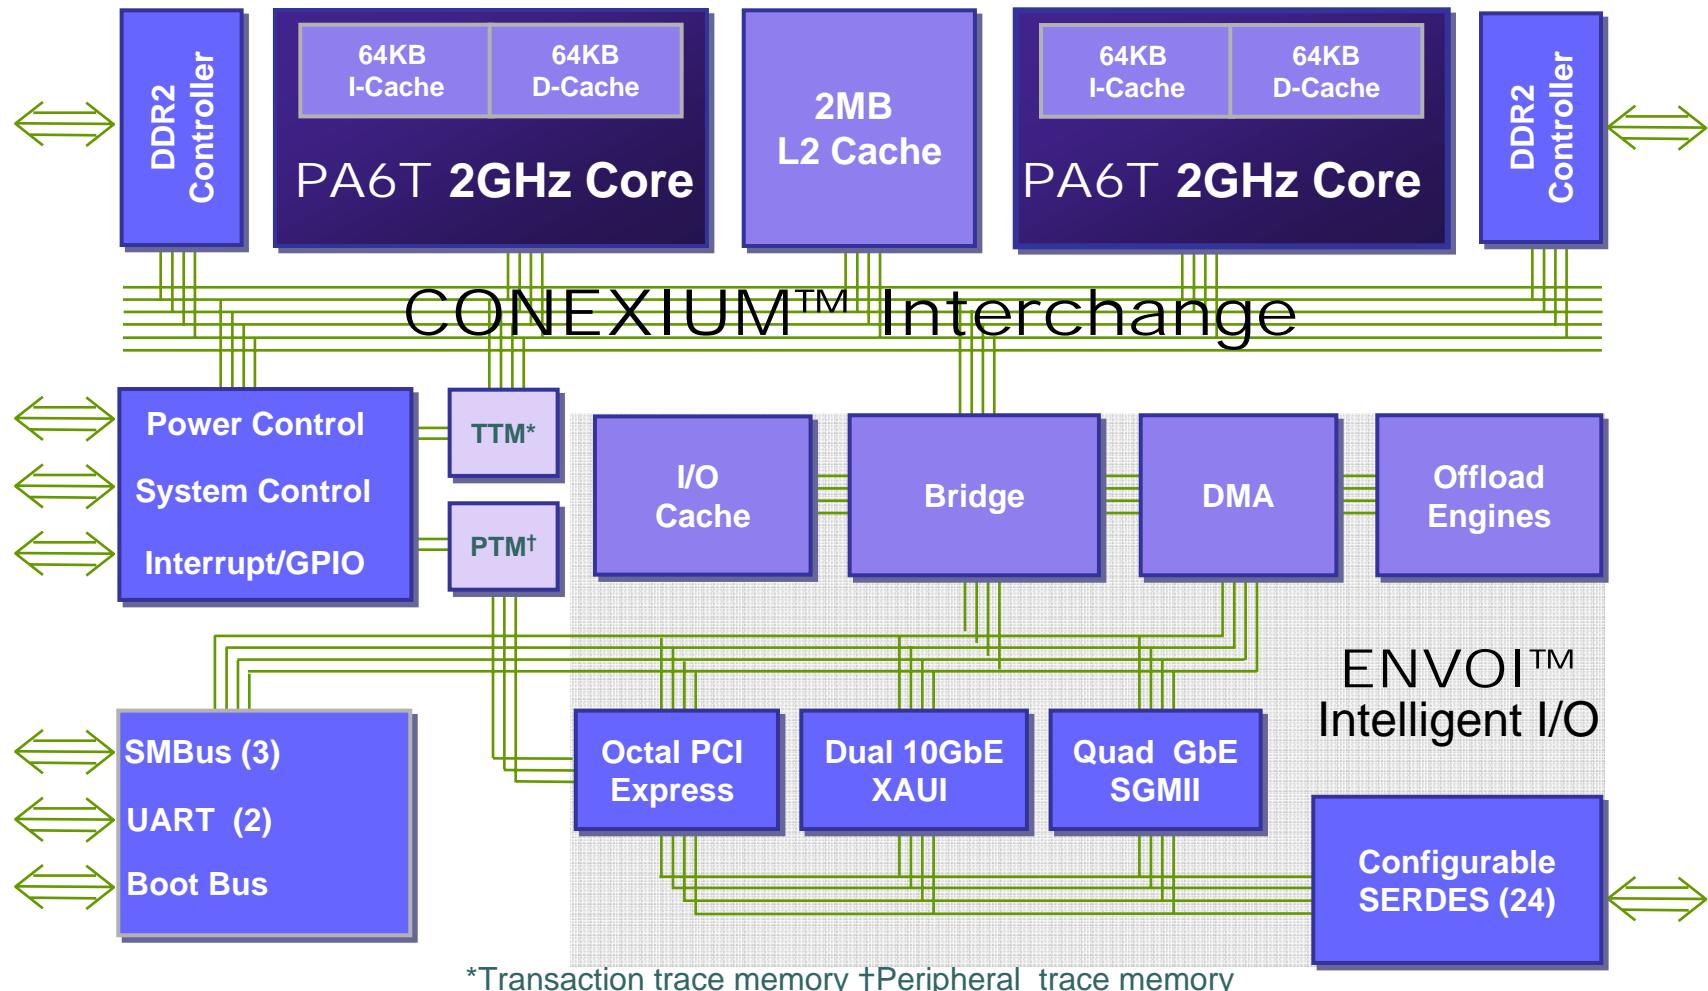
- ▷ Saves interface power

## ► Management

- ▷ Voltage/frequency scaling
- ▷ Multiple power planes for optimal voltage selection per region
- ▷ Clock gating to reduce power of idle circuits
- ▷ Active and pre-charge standby modes in DRAM array
- ▷ PCIe power saving modes
- ▷ Nap and Sleep modes for CPU

# PWRficient Components

# PWRficient 1682M Block Diagram



# Power Efficient

Power is first-order design principle

7W >15,000 gated  
7W Clocks on chip

1.5W separate power  
rails for core, IO

1W  
PADs

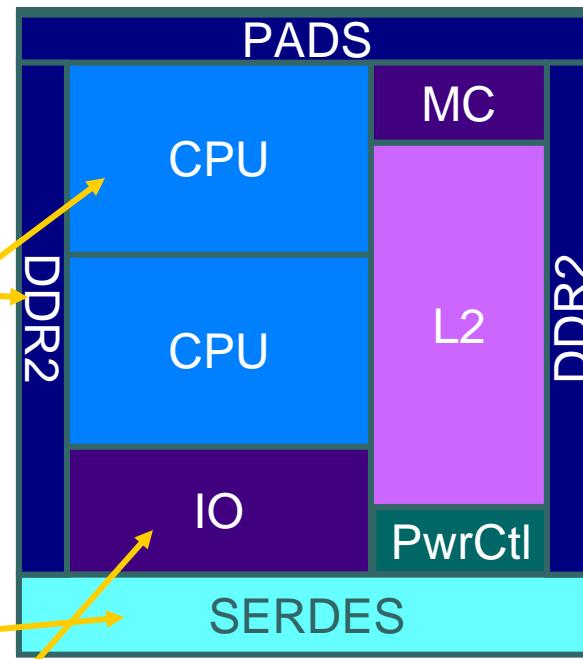
1W

5W

1W

2W

Turn off unused I/Os



MC optimizes  
DRAM power

L2 cache and I/O  
coherent with core  
off

Dynamic Power control hardware and  
software voltage/frequency management

# PA6T Core

- ▶ Fully compliant Power implementation
  - ▷ Power 2.04 architecture spec
  - ▷ Full FPU and VMX SIMD support
  - ▷ 32- and 64-bit
  - ▷ Big- and little-endian modes
    - ▷ Bi-endian per process
- ▶ Super-scalar, out-of-order design
  - ▷ Quad fetch, triple issue
  - ▷ 64-entry scheduler
  - ▷ Strongly ordered memory model
- ▶ Hypervisor and virtualization support
- ▶ High performance memory hierarchy @ 2Ghz
  - ▷ L1 data—32GB read or write
  - ▷ L2 data—16GB read plus 16GB write
  - ▷ Memory DDR2-1066 —16GB read or write
  - ▷ 16 transactions in flight
- ▶ CONEXIUM Interchange
  - ▷ 1G transactions per second
  - ▷ 64GB/sec data peak
- ▶ Extensive power management capabilities

# Processor Pipeline

## Multi-Issue, Out-of-Order, Superpipelined

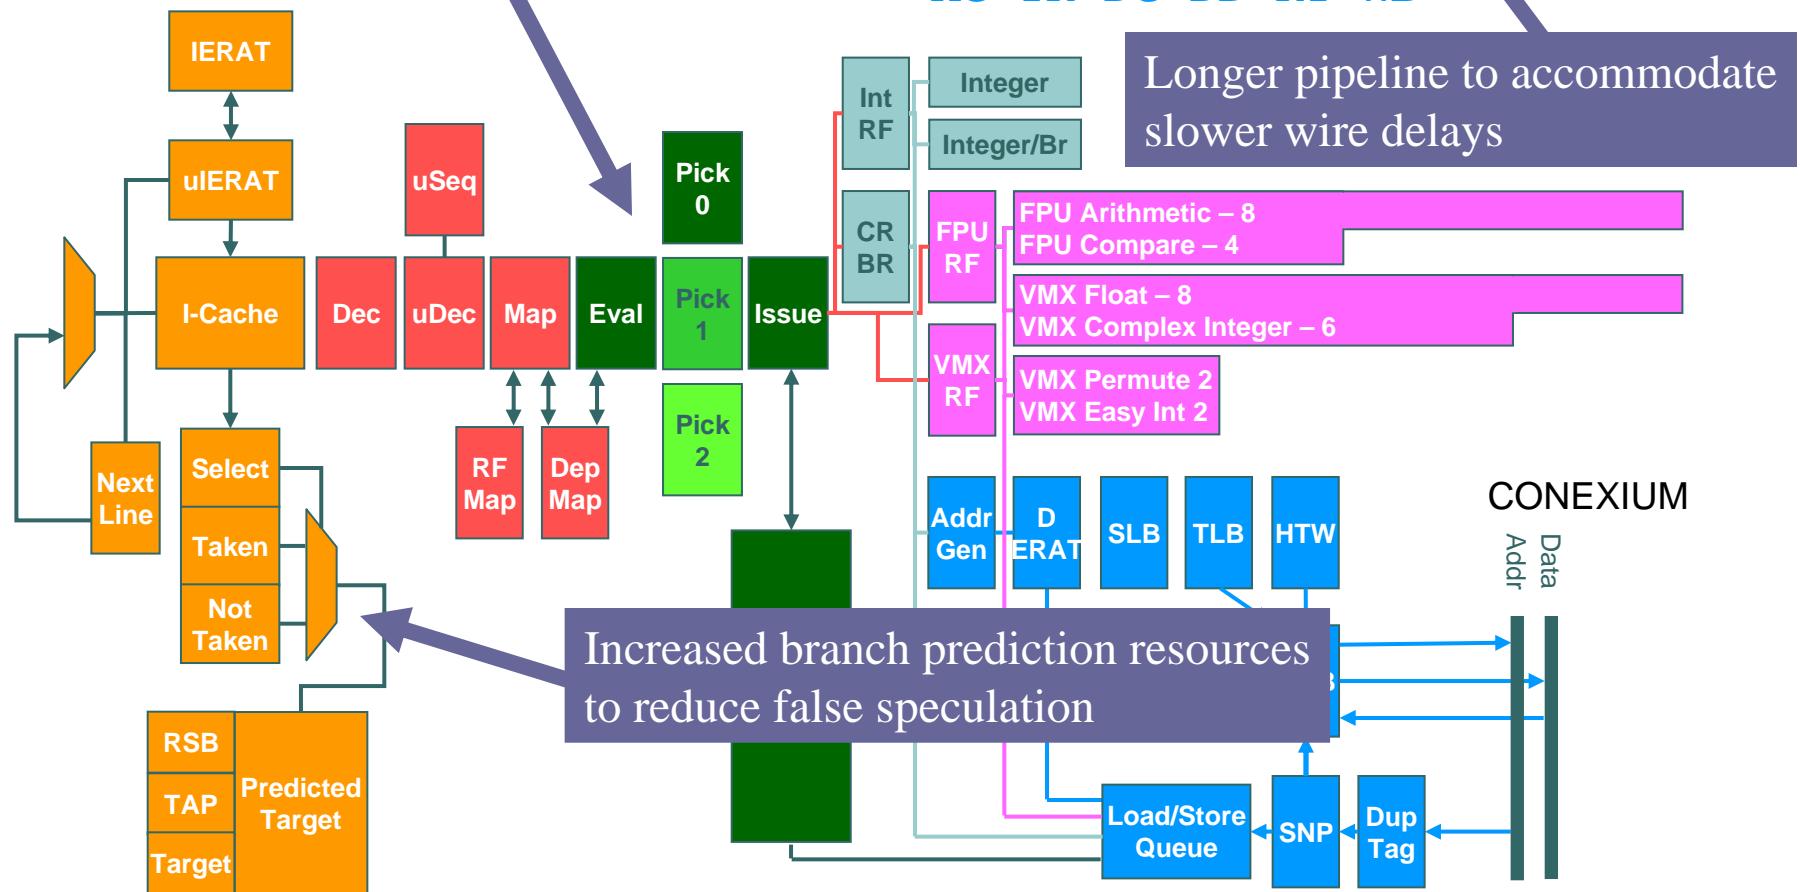
Reduced instruction issue width to save power ~  $N^2$  to  $N^3$

7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19

IF IF IC DE UC MP SE SP IS RR AL AD RT WB

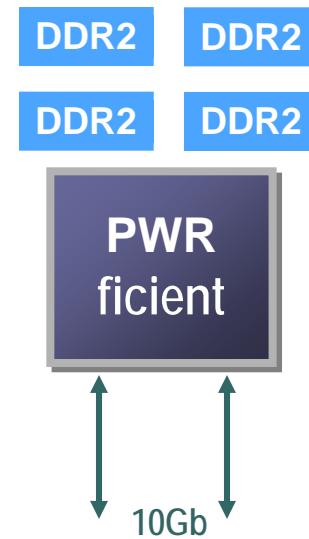
AG TR DC DD RT WB

Longer pipeline to accommodate slower wire delays



# DDR2 Memory Controller

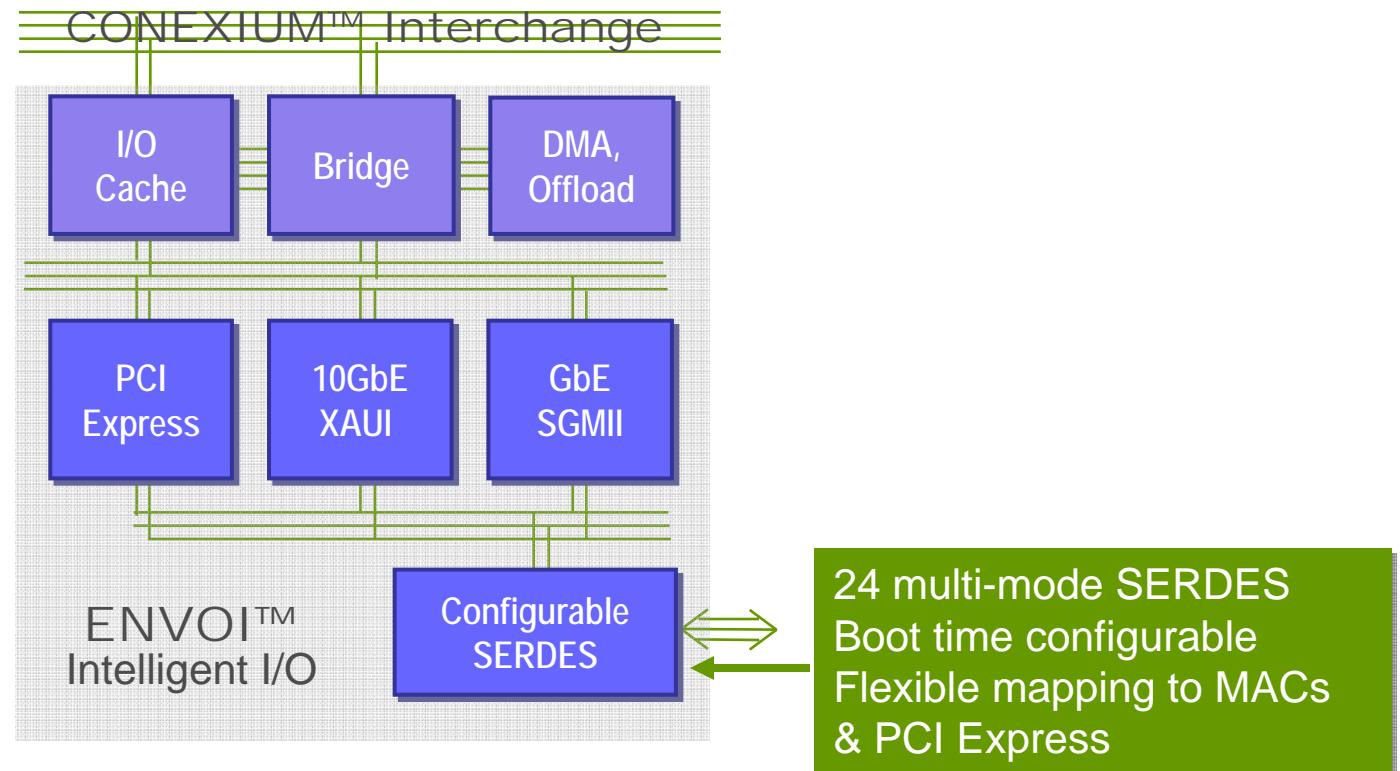
- ▶ 5 Watts for active DRAM rank
- ▶ 1 Watt for standby DRAM rank
- ▶ DRAM is a significant percentage of system power
  - ▶ SOC 13 Watt typical
  - ▶ 2 x 2 DIMM memory is ~12 Watts
- ▶ DRAM controller manages power
  - ▶ Extensive clock gating within the controller
  - ▶ Use of DDR2 power saving modes
    - ▶ Precharge standby
    - ▶ Active standby
  - ▶ Sorts requests to maximize power saving



# ENVOI — Intelligent I/O

- ▶ Interoperability between PCI Express, packets, DMA, and memory

- ▶ Centralized DMA model
- ▶ Shared resources saves area and power, which enables integration
- ▶ Idle functions are clock gated or powered off



# Performance and Power

# High-Performance At Low-Power Across A Range Of Metrics

PWRficient 1682M provides mainstream performance at low power

## General-purpose computing

- ▶ SPECint™ >1000 per core\*

## Floating-point performance

- ▶ SPECfp™ >2000 per core\*
- ▶ Imaging
  - ▷ FFT 24 GFlops/sec (total)\*

## System bandwidth

- ▶ Memory-to-memory copy
  - ▷ 10 Gigabytes/sec\*
- ▶ High-speed serial I/O
  - ▷ 24 SERDES for total I/O bandwidth up to 104Gbps peak

## Application offloads

- ▶ TCP/IP termination
  - ▷ > 20Gbps†
- ▶ Encryption
  - ▷ 10Gbps bulk IPSec and SSL encryption†
  - ▷ 3,000 public-key handshakes/sec in software\*
- ▶ Storage
  - ▷ 2.0GB/s RAID5 (data+parity)†

\*Estimated max sustained performance at 2GHz

†Estimated peak performance at 2GHz

# Power Summary

	Typical Power	Worst Case Power	Comments
PA6T Core @ 2Ghz	4W	7W	Core-only power
Platform, dual-core @ 2GHz	13W	25W	Full SOC, everything on
Platform, nap mode	1W	2W	Coherent L2 & IO

# Power-Area Efficiency Index

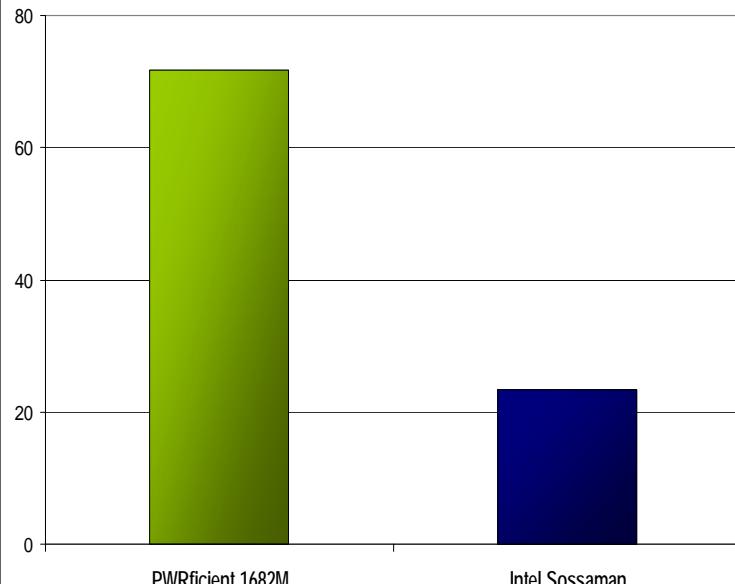
- ▶ Customers care about performance, power, and board area
- ▶ P.A. Semi uses a Power-Area Index as a measure of computing efficiency
  - Captures both power and space efficiency

**PA**  
Power Area  
Index

=

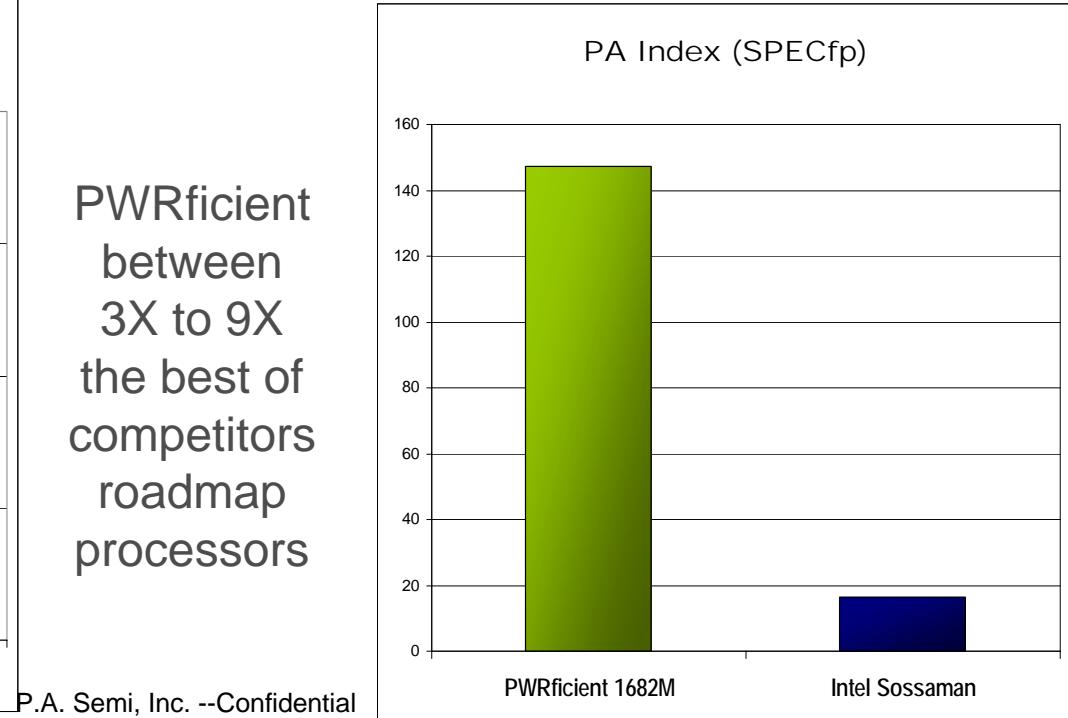
$$\frac{1000 * \text{SPEC performance}}{\text{Max. Power (W)} * \text{Platform Board Area (sq mm)}}$$

PA Index (SPECint)



PWRficient  
between  
3X to 9X  
the best of  
competitors  
roadmap  
processors

PA Index (SPECfp)



# Supercomputer 500 Case Study



## 32 Teraflop Supercomputer

PWRficient dual-core cluster versus

- Opteron dual-core cluster
- Xeon dual-core cluster

### Space Efficiency

#### # Node Clusters

- ▶ 4,000-node PWRficient cluster
- ▶ 8,000-node AMD Opteron cluster
- ▶ 5,333-node Intel Xeon cluster

### Cost Efficiency

#### Platform Silicon\* Cost

- ▶ \$1.4M for PWRficient 2GHz
- ▶ \$4.6M for AMD 2.2GHz Opteron
- ▶ \$3.3M for Intel 3.2GHz Xeon

### Power Efficiency

#### Power Cost Over 3 years

- ▶ \$360K PWRficient cluster
- ▶ \$3.5M for AMD Opteron cluster
- ▶ \$3M for Intel Xeon cluster

\*Processor, plus northbridge, southbridge, 10G MAC, if required

# Summary

- ▶ New Technology and SERDES I/O enable new possibilities
- ▶ Optimize performance, power, latency, density
- ▶ From-scratch design maximizes the gains
- ▶ Interesting convergence of needs across a wide range of design points
  - ▷ Networking, telecom, servers, imaging
- ▶ High performance and low power enable wide range of applications